The new challenges of High Performance Computing

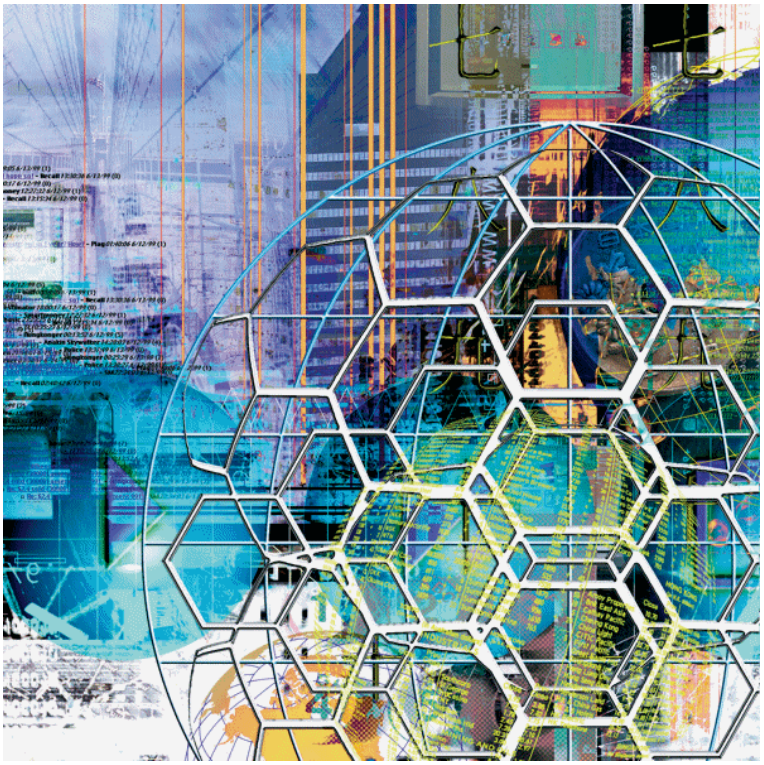

Open and

innovative solutions

for High Performance

Computing

www.bull.com

Bull

“The Bull solutions in the field of High Performance Computing are the fruit of a long history of expertise in large server architectures and open software. These solutions are particularly relevant considering new production conditions. The introduction of technologies based on standard components, and in particular Intel 64-bit processors, represents a real breakthrough in terms of price/performance and investment protection. Bull wants its customers to take advantage of this breakthrough, so that public and private research centres can have much wider access to best-of-breed solutions in High Performance Computing”

Michel Guillemet
Research and Development Director, Bull

A revolution is under way

New uses

High Performance Computing (HPC) is a major issue for society. Originally intended for a limited number of fields, such as nuclear physics and meteorology, HPC is now meeting increasingly wider needs. Today's society requires better knowledge and understanding of physical phenomena, the processes of living and economic activity, and better ways to design and produce new products. As a consequence, public and private research centres must have access to increasingly higher data processing power to achieve these objectives.

New technologies

Traditionally, High Performance Computing was based on extremely expensive solutions using specific processors and software applications, which considerably limited its use. Today, with standard processors and Open Source software, the strong need for High Performance Computing can be satisfied. These new technologies make it possible to offer the highest performance with a cost of ownership similar to that of standard products.

The impact of data

A new paradox: if the power provided by processors is a determining factor for High Performance Computing, then data storage and management solutions are also fundamental. As a matter of fact, as the volume of the data to be processed becomes increasingly higher, data access times must be very fast and this data must be better integrated into the IT system.

Bull's commitments

In this new context, Bull decided to be a major contributor by leveraging its expertise in large server architectures based on standard components, in middleware and Open Source, and in data management and the integration of complex systems. Bull, today, is addressing this sector with innovative approaches and products that will have a major impact and will bring progress in a broad field of activities.

Three questions to Professor William Jalby

PRiSM Research Laboratory in Computer Science
University of Versailles

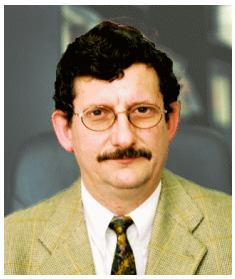

Professor Jalby, in your opinion, what are the fundamental contributions of Intel® Itanium® 2 to the evolution of high-performance processors?

The IA-64 instruction set was designed to take into account both recent and future evolutions in microprocessor architecture. The analysis carried out by Intel indicates that we are close to the limits of "out-of-order" technologies, which let the hardware discover and exploit the potential sources of parallelism. Intel, therefore, decided to make a qualitative jump by carrying out the search for parallelism at the program preparation level. The savings thus made were used to ensure the exceptional performance of the Itanium® family. I am referring more particularly to the massive hardware resources (functional units, registers, memory hierarchy...), predication and speculation, the register stack engine, fused multiply/add unit.

The result? After a first generation of processors that did not measure up to expectations, the Itanium® family is among the top few which until now included only the proprietary chips of IBM, Compaq and HP. To our knowledge, the IPF(*) roadmap indicates a continuous growth in performance and functionality that could well outstrip certain competitive architectures.

"We particularly appreciate the density and performance of NovaScale, its low NUMA factor, and its excellent price/performance ratio"

Another innovation in Intel's approach is in its support for system manufacturers. Unlike x86 processors, Intel® Itanium® 2 comes with companion chipsets that help OEMs design large multiprocessor systems. At this point, I am thinking of course of Bull's NovaScale™ server, which has been used by our students and researchers. We particularly appreciate the density and performance of NovaScale™, its low NUMA factor compared with architectures that we used in the past, and its excellent price/performance ratio. Other manufacturers are proposing or will propose top-of-the-range systems based on the Itanium® family, which justifies our investments - both hardware and intellectual - in this architecture.

(*) Itanium® Processor Family

Can these systems based on Intel® Itanium® 2 address both the scientific and business worlds equally well? For what kind of applications?

I think that with the availability of processors such as Intel® Itanium® 2 there will be no more specialised scientific processors or systems. Typically, the processors of the x86 family poorly addressed the performance requirements of scientific workloads, in terms of floating point processing, memory throughput rates, and vector computation. Itanium®, while maintaining its qualities for business applications, has made an impressive entry into the scientific and technical world, even if the progress achieved in the three above-mentioned characteristics is still unequal.

"Itanium, while maintaining its quality for business-critical applications, has made an impressive entry into the scientific and technical world"

In fact, the work achieved on floating-point processing performance is rather exceptional. One just has to look at the SPECfp benchmark to be convinced. The memory throughput, even if it has jumped compared to other Intel processors, remains a little low for applications that need a "working set" size larger than the cache. This is a problem common to all SMP server architectures based on a system bus. The only exception is the vectorial architectures, which are very expensive. Moreover, there are now only two manufacturers who still sell vectorial machines…

Most of the work of my team concerns the optimisation of vectorial workloads on scalar parallel systems. We think we have found the way to greatly improve the performance of these applications. This is another reason why I appreciate Itanium®, because it gives back to software developers the ability to optimise programs in depth - a possibility which so far was largely confiscated by the "super-scalar out-of-order" approach.

In other words, I would say: yes, these systems are appropriate for both scientific and technical applications, with the exception of a few areas that still remain dependent on vectorial architectures.

Three questions to Professor William Jalby (cont.)

Which software environments can make the most of Intel® Itanium® 2 performance?

The decision made by Intel to discontinue a very long line of processors that were 100%-compatible at the binary level has much less impact on the scientific and technical markets than on the business world. First of all, this segment is less dependent on the availability of applications provided by Independent Software Vendors (ISVs), because they are often specific to each user. Once the appropriate development tools are available, users will be able to implement their own Itanium®-based system environments.

"The adoption, since quite a few years now, of Linux and Open Source as the preferred software environment has been a decisive element which makes things simpler in the scientific field"

The other decisive element that makes things simpler in the scientific field has been the adoption, since quite a few years now, of Linux® and Open Source as the preferred software environment. This market has now reached the critical mass necessary to attract developers and software publishers, without forgetting the strike force of thousands of volunteer developers, which no proprietary environment could equal.

The principal challenge that manufacturers and Linux® developers have yet to face is the evolution of the operating system to effectively manage the specifics of large multiprocessor systems. Until now, the industry has circumvented the problem by building clusters of small systems (two or four processors). To make the most of large multiprocessors such as Bull NovaScale™, it is advisable to further extend the improvements introduced by projects like Atlas.

A great potential for progress

High Performance Computing is an important motor for progress in the world today. In practically all research sectors, one finds challenges for the future that can be solved only through advances in this field. In industry, improving computing capacity without increasing cost makes it possible to design new products that would otherwise have been too complex or too expensive to create, and increases the productivity.

Here are just a few of the many examples that show the importance of High Performance Computing and its impact on society.

Improved weather forecasts

Weather forecasting was one of the first large-scale uses of High Performance Computing. But considerable progress still needs to be made to help industry sectors, such as agriculture and transport and to provide better advance warning of natural disasters.

This progress is now possible, thanks to the increasing amounts of data becoming available to re-adjust models, and to the availability of higher computing power that will enable the use of more precise models.

▪ On the first of these points, data collection is currently taking a great leap forward, most notably with new weather satellites providing large quantities of high-precision information.

▪ On the second point, Japanese researchers have shown, with the "Earth Simulator", the kind of progress that can be made with more significant computing power. This IT system is powerful enough to provide simulations with a resolution of 10 km, compared to 100 km in the past. This improved precision makes it possible to forecast the formation of tornadoes, which did not appear in former simulations.

In this way, a significant improvement in computing power brings concrete benefits in the quality of the forecasts.

Brain scanning for preventing and treating diseases

In this field, the introduction of new technologies for providing brain images, such as Positron Emission Tomography (PET), Magnetic Resonance Imaging (MRI), and the magneto-encephalography (MEG), have introduced a new need for High Performance Computing. Detecting temporal evolution, providing appropriate visualisations and analysing the data resulting from these imaging techniques require powerful means of calculation.

The availability, at reasonable price, of High Performance Computing systems would, for example, make easier to locate the modifications of chemical substances and thus draw up a thorough cartography of the functions of the brain. This could be used to improve our understanding of pathologies like epilepsy and Parkinson's disease and, then envisage solutions for developing personalised treatments thanks to the real-time visualisation of the effects of therapy.

Such progress would also lead to a better understanding of how the brain functions, which is certainly one of the most enthralling challenges of the 21st century.

Virtual prototypes to reduce costs, shorten implementation times, and improve product quality

Several industrial sectors, such as aeronautics and automobiles, invest in High Performance Computing to improve design methods for their products. In particular, these industries have evolved to a virtual prototype methodology where cars and planes are now completely simulated before being built.

By removing the trial-and-error cycle on real prototypes, the recourse to simulation makes it possible to shorten implementation times. Simulation also makes it possible to improve productivity and reduce costs - real prototypes being generally more expensive than substitute calculations.

By adopting this method of virtual prototypes, companies benefit from an increase in the performance of their simulations at equivalent or lower cost. Any advances in High Performance Computing can thus have an immediate impact on productivity.

Greater calculation capacity also means increased safety by allowing more complete simulation of the safety-related factors at a reasonable cost for manufacturers.

Three requirements to make High Performance Computing more accessible

Making High Performance Computing more accessible to research and industry is a true challenge for the IT community. The three essential requirements to take up this challenge are the association of lower cost with a strong increase in the available power, control of data, and guaranteed long-lasting solutions.

Improved price / performance ratio

The improvement in the price/performance ratio is necessary for the research teams to respond to the scientific and technical challenges of today's world. Once this requirement is fulfilled, it will be possible to deal with more complex problems, increase the precision of models, and combine several models and progress simulations along with the computing results obtained.

In manufacturing, the fall in the unit cost of computing power will make it possible to improve productivity. This will result in more powerful and innovative products that were not previously possible for cost and profitability reasons.

Efficient data management and integration into the IT system

But having computing power can be insufficient without high-performance data management solutions. In fact, these applications use more and more data and generate more and more results. The productivity of scientific teams is often directly related to the capacity of the systems to manage these huge volumes of data.

To be able to effectively store hierarchically-arranged data on various levels of a storage system is essential for the effectiveness and efficiency of an IT system. This requires great know-how in the balancing of computing and storage architectures. Even the handling of data is a significant factor in the correct operation of computer centres. To know how to manage data according to multiple characteristics is crucial for its effectiveness and productivity.

Durable solutions sharable between laboratories

Durability is also a major requirement. The applications have a very long life span and were developed at the expense of great effort. It is not possible to re-examine their design with each change of an IT system. To minimise modifications, there is a strong demand for solutions based on standard software.

The interdisciplinary character of scientific computing reinforces this need. The interactions between researchers, algorithmicians or numericians, and data processing specialists are numerous. Thus, open and standard solutions make it possible to re-use existing applications and to facilitate their evolution, thereby increasing the productivity of the teams.

Proud to contribute to scientific research and industrial innovation

Gérard Roucairol, Chief Scientist, Bull

With the introduction of NovaScale™, the new large open servers based on the FAME[1] architecture and Intel® Itanium® 2 processors, Bull has once again teamed up with universities, public research centres, and the engineering departments of the automobile, aeronautics, pharmaceutical and other industries.

In fact, the FAME architecture is particularly appropriate for applications requiring significant computing power. The Intel® Itanium® 2 processors integrate floating-point calculation units and have multi-media instructions. Moreover, they benefit from innovations such as parallelism between the instructions of a program.

The SMP[2] architecture with shared memory is the most effective solution for multithreading applications (written for several processes being executed in parallel). However, the needs for the most demanding applications in terms of computing power exceed the capacities of the largest SMPs available. This led the industry to combine the multiprocessor programming model with shared memory programming by "message passing". For these applications, the Bull NovaScale™ servers can be interconnected by a low-latency high-bandwidth link to build clusters responding to the highest requirements.

Complementing the quality of the NovaScale™ platform, Bull's European presence and the ready availability of Bull experts are major assets to users seeking to implement solutions that precisely meet their needs. A close relationship with the R&D engineers can be critical for users when adjusting system parameters and adapting them to the constraints of the applications, in order to maximise system performance. This proximity with the R&D teams makes it possible for Bull to build a true partner relationship with laboratories and research centres.

Bull is proud to be able to contribute to scientific research and industrial innovation.

[1] Flexible Architecture for Multiple Environments
[2] Symmetric Multi Processors

Bull's High Performance Computing solutions

Bull has developed a whole range of solutions that provide significant innovations to meet the needs of High Performance Computing. These solutions benefit from the experience accumulated by Bull over many years in the management of large system architectures, operating systems, software environments and administration.

Build powerful servers around standard components and based on an innovative architecture

To build its High Performance Computing solutions, Bull has given priority to the price/performance ratio, while offering the power and memory necessary for the most demanding applications.

In order to achieve this goal, it is imperative to build solutions based on standard components which, because of their large distribution, also have a low price. In addition to the learning curve effect related to mass production, the design and marketing effort of these types of components can be amortised through a large distribution, giving them an indisputable advantage in terms of component prices over a low volume distribution.

Figure 1
Bull NovaScale™ Servers

Associating these standard components to produce powerful, robust and scalable servers is the objective of the FAME architecture that Bull designed and developed for NovaScale™, its new generation of large open servers.

Intel® Itanium® 2: the most powerful as of today

The choice of IPF (Itanium® Processor Family) rises from Bull's strategy to offer the best price/performance ratio. The Intel® Itanium® 2 processor is, today, the most powerful processor based on a significant number of benchmarks. The new "Madison" version will further increase this pre-eminence of IPF on the 64-bit processor market. This performance is due to the technological innovation developed by Intel with its EPIC (Explicitly Parallel Instruction Computing) architecture, which makes it possible for applications to exploit the parallel resources of the processor. This new technology still has a strong improvement potential that guarantees an IPF increase in performance over the course of time - one that will be largely higher than that of older competitors' architectures that have already exhausted to a great extent their optimisation possibilities.

Intel® Itanium®: a guarantee of longevity

The strategic commitment of Intel on 64-bit processors makes it possible to guarantee longevity with this family of processors. The processor roadmap is clear and compatibility guaranteed between generations ensures the manufacturers, software publishers and users that their investments on these processors could be made profitable over a long period of time.

Bull collaborates with Intel on IPF since 1997

Bull, having identified very early the potential of the IPF, established a collaboration with Intel in 1997, which enabled them to design and develop a range of servers based on this family of processors. Thanks to this co-operation, Bull has a fine knowledge of the possibilities offered by the Itanium® processors and could make the most relevant choices to build its line of NovaScale™ servers.

FAME, the innovative architecture of Bull NovaScale™ servers

The NovaScale servers benefit from the FAME architecture developed by Bull to achieve scalable multiprocessor systems with shared memory providing an exceptional price/performance ratio. The FAME architecture is based on the implementation by Bull of very high-speed interconnections and on the development of a very sophisticated chip: the FSS (FAME Scalability Switch) which ensures that each processor has access to the I/O and has a coherent vision of the global memory. This memory is formed by uniting the memories associated with each quadriprocessor board.
According to the size of the server, the number of FSS used increases to ensure a balance of communications in the system.

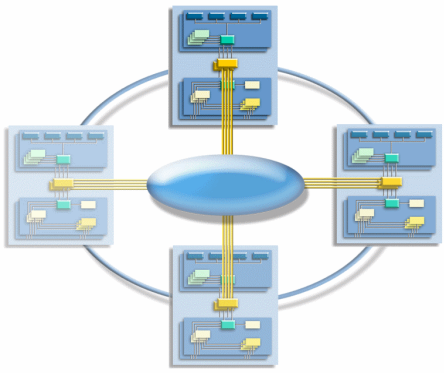

Figure 2
The FAME architecture of large Bull NovaScale™ servers

The first Bull NovaScale™ servers benefiting from this technology were introduced in March, 2003 with models accommodating up to 16 Itanium® 2 processors. A model with 32 processors will be soon available.
The Bull NovaScale™ servers can be integrated, via high-speed networks, to larger computing infrastructures, either locally (clusters), or over long distances (grids).

An exceptional price / performance ratio

Thanks to these choices and to the innovative architecture of the NovaScale™ range, Bull is today at the head, in front of the whole of its competitors on the price/performance ratio for its range of 64-bit servers. A study based on the SPECfp_rate benchmark and the list prices proposed by the manufacturers show an advantage for the Bull NovaScale™ offer as well for the 4, 8 and 16-way servers.

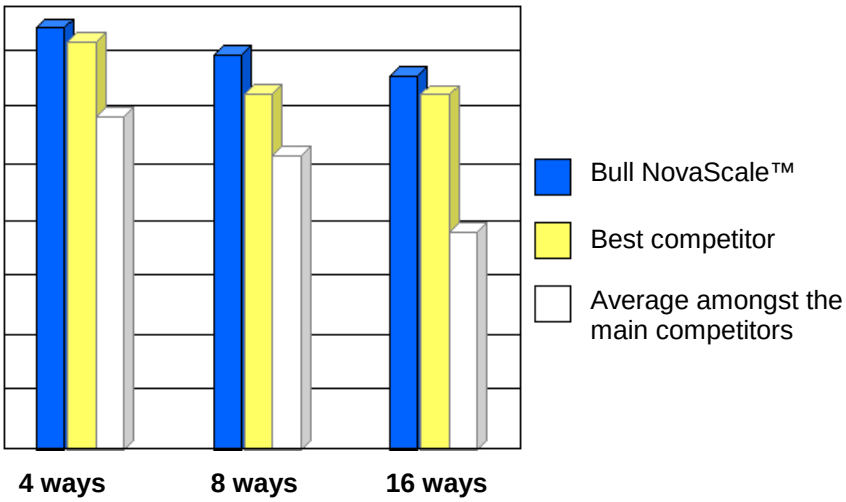

Figure 3
Servers price/performance ratio (May 2003)

This graph proves the relevance of the Bull approach based on the use of standard components to build a High Performance Computing offer which meets the needs for this sector – that is having solutions offering the best price/performance ratio.

SMP cluster for highly powered configurations

The Bull NovaScale™ servers range has been designed to be configured in clusters to meet the needs of applications requiring high computing power. Bull developed, with its partners, clustering solutions offering both excellent communication performance between servers and a complete management environment, which makes it possible to ensure the operation of all of the servers.

Since the FAME architecture offers a linear price progression between 4, 8 or 16 ways, it is possible for the same level of power to choose between a cluster of a great number of compact servers with a limited number of processors and a cluster composed of a reduced number of powerful servers with a significant number of processors. For example, a 16-way cluster node is not more expensive than four interconnected nodes of 4-way servers.

Bull makes it possible for customers to choose the granularity of the nodes of their High Performance Computing system according to the needs of their applications and not according to financial considerations. This is particularly interesting for users having applications that require nodes with a significant memory size and significant local power.
This also makes it possible to have powerful clusters with a relatively limited number of nodes and thus less complex to manage.

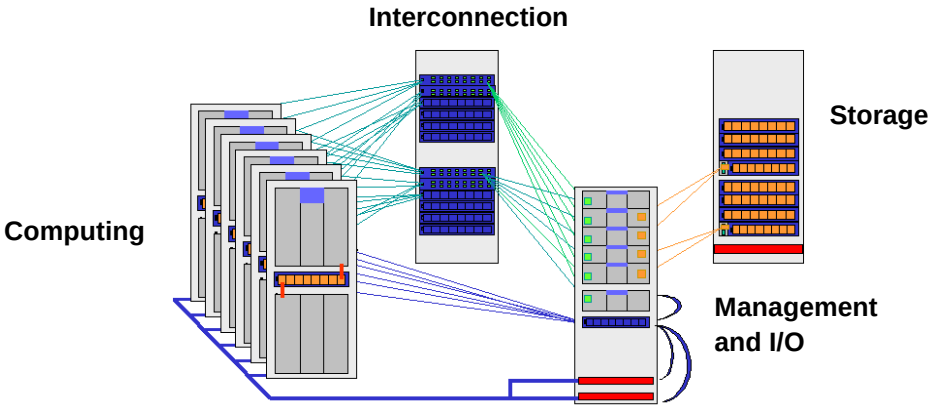

Figure 4
FAME architecture server cluster

To operate these clusters, Bull developed a complete software environment making it possible to control the whole configuration as a single system. It integrates software deployment functions, resource control and monitoring all from a centralised management console. Users have a traditional job scheduler interface ensuring the execution of their work on the clusters. All of these functions use either Open Source software or standard market software, and are integrated and optimised to allow for easy and powerful operation on the Bull NovaScale™ clusters.

A long experience in data architectures

Bull has great expertise in the control of complex IT infrastructures, combining computing power and data management capability. Bull has, for many years, provided solutions offering high performance and availability, and has strong competencies in the choice of architectures, which make it possible to optimise data management.
This experience is a significant asset in the field of High Performance Computing, which is increasingly faced with problems in balancing capacities of calculation, I/O throughput and the hierarchical organisation of storage. Bull proposes for its clusters, a global file system software and invests in solutions that are able to take into account the increasingly high requirements of large computer centres, such as the Lustre File System.

An open and standard environment

For the construction of its High Performance Computing offer, Bull chose to use standard open software, from the world of the open software (Open Source), or from independent publishers. Application code can thus have a significant life span, and the environment's hardware changes can result in reduced cost for the teams of researchers, numericians and IT specialists since they can preserve the same software environment.
For the community of High Performance Computing, which already has a strong tradition of co-operative developments, Open Source allows the mutual investment necessary to develop applications and for the contribution of new functionalities. Certain independent publishers also propose powerful software which become de facto standards. This software is available on the great majority of the platforms. The users have the assurance that the products will be maintained in the state of art and will provide real protection of their investments.
On the other hand, specific manufacturer's software cannot provide this guarantee. They offer interesting functionalities but are dependent on platforms and run the risk of seeing their competitiveness decline.

Optimised software on Bull NovaScale™ servers, a strong commitment on Linux®

The solutions that Bull selected make it possible for users to have a complete range of innovative, open and long-lasting software. Bull optimises these solutions on its servers, while preserving their open and standard character. Bull strongly invested to make it possible for Linux® to support the functionalities required in High Performance Computing solutions.
Bull is an active member of the Atlas Open Source project, which gathers the principal actors in the Linux® world to allow the operation of this system on servers built on Intel® Itanium® processors. The Bull teams also take part in projects like Linux® Test Project (guarantee of Linux® robustness) or Linux Scalability Effort (support of large SMP and NUMA servers) which today make it possible to have an implementation of Linux® that meets the needs of High Performance Computing.

A complete offer to develop and implement applications

Bull's scientific computing offer supports the whole range of parallel programming models, in particular OpenMP and the PVM and MPI libraries. Bull proposes a standard version of MPI, optimised to the FAME architecture and the power of Bull NovaScale™ servers. In order to exploit the potential of the EPIC architecture on Itanium® processors which enables the execution of several instructions in parallel, parallelism must be detected and processed at the compiler level. For these reasons, Bull proposes the Intel (C/C++ and FORTRAN) compilers since Intel teams have the competence and architecture knowledge on these processors to provide powerful products.

Complementary to the compilers, the scientific libraries make it possible to make the most of the Itanium® processors performance. Bull offers the Intel optimised MKL scientific library, the IPP (Integrated Performance Primitives) library, as well as the Open Source libraries, which provide a large variety of functionalities.
Bull's scientific environment also has powerful application debugging and optimisation tools. Based on market proven software, such as Vampir, TotalView®, or Open Source solutions integrated with the Bull NovaScale™ servers, these tools give the users a number of different means to make the most of the available power.

Dedicated scientific and technical Competence Centres

Bull's HPC solutions are supported by experts who are able to help users benefit from the available innovations. Traditional Bull competencies on large systems were reinforced by its work in the world of Open Source and by its co-operation with scientific and technical computing actors. Two Competence Centres, around Open Source and HPC, were set up to help users fully benefit from the solutions proposed by Bull.

Application optimisation on Bull NovaScale™ servers including Open Source tools and storage solutions

Even if automatic tools constantly progress, the complexity of processor architectures makes it still necessary to optimise to obtain the best possible results with the applications. Through its co-operation with university research teams and based on its own experience, Bull offers high level services so that the user's applications can fully exploit all the power available while drawing on the specificities of the Itanium® processor family and its architecture.

Bull's commitment in the Open Source community, its knowledge of software tools, its experience in management solutions including high data volumes, and its capacity to deploy and support large projects are major assets to implement powerful, reliable and long lasting solutions.

A world of standard and open solutions for High Performance Computing

With Bull’s High Performance Computing solutions, users have a powerful offer using the industry's best market standards, at the best price and with the greatest possible longevity.

The adoption of the Intel® Itanium® processor family is a major asset, ensuring the best level of performance today and in the long term. The development of an innovative architecture for the Bull NovaScale™ servers is a perfect solution to meet the user's need for powerful servers that can be easily managed and integrated into a wide number of configurations, clusters and grids.

With an open and complete software environment, users will be able to benefit from their developments over a long period of time and, if they wish, share these developments within their community.

Bull’s Competence Centres and the co-operation developed with recognised research centres make it possible to bring a very high level of expertise to users.

Bull, a European manufacturer, is proud to be able to help today’s world progress in the understanding and control of environmental issues, to acquire new knowledge and to develop innovative and powerful products.

Bull S.A. - 68 Route de Versailles - 78430 - Louveciennes – France
UK: Maxted Road, Hemel Hempstead, Herts HP2 7DZ
USA: 300 Concord Road, Billerica, MA 01821

www.bull.com

Version VE 01